

algorithms for statistical profiling are not incorporated-by-reference into *Emerald 1997*, it would have been obvious to one of skill in the art to combine *Emerald 1997* with prior NIDES-related publications, since *Emerald 1997* states that the NIDES algorithms form the basis for the statistical profiling disclosed. One of skill would have been motivated to further investigate these specific algorithms in order to improve the performance of the system.

258. In my opinion, as more fully reflected in Exhibit H to my report, the system disclosed in *Emerald 1997* satisfied every limitation of each of the asserted claims of the patents-in-suit (except '338 claim 7) and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the inventions claimed in the patents-in-suit. Consequently, the noted claims of the patents-in-suit were not novel when filed.

259. In addition, as reflected in Exhibit H, the remaining indicated claim is invalid as obvious, for the reasons detailed above.

a. CMAD and Conceptual Overview

260. In addition to the *Emerald 1997* paper, SRI also published two shorter earlier descriptions of the EMERALD system. See P. Neumann, P. Porras and A. Valdes, "Analysis and Response for Intrusion Detection in Large Networks," Summary for CMAD Workshop, Monterey, 12-14 November 1996 ("Emerald CMAD") and P. Porras and P. Neumann, "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances Conceptual Overview," December 18, 1996 ("Emerald – Conceptual Overview"). While neither of these references is as detailed as *Emerald 1997*, both also anticipate or render obvious many of the asserted claims of the patents-in-suit for the reasons stated above.

that the techniques described in this paper may be used for the "profile engine" which "can profile network activity via one or more variables called measures." '338 col. 5:43-50.

b. Live Traffic Analysis and Conceptual Design 1997

261. SRI also published two additional papers regarding the EMERALD system: P. Porras and A. Valdes, "Live Traffic Analysis of TCP/IP Gateways," <http://www.sdl.sri.com/projects/emerald/live-traffic.html>, Internet Society's Networks and Distributed Systems Security Symposium, Nov. 10, 1997 ("*Live Traffic Analysis*") and P. Porras and P. Neumann, "Conceptual Design and Planning for EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances," Version 1.2 20 May 1997 ("Emerald - Conceptual Design 1997"). I understand that there is a debate over whether or not these documents were publicly available prior to November 9, 1997. If public, *Live Traffic Analysis* anticipates all of the asserted claims of all of the patents-in-suit, except for '338 claim 25, which is rendered obvious, as indicated in Exhibit L. If public, Emerald – Conceptual Design 1997 anticipates or renders obvious all of the asserted claims of the patents-in-suit, for the reasons discussed above and those noted on Exhibit K.

c. Emerald 1997, Intrusive Activity 1991, NIDES 1994

262. *Emerald 1997* also references two additional publications in both the text of the paper itself, and in the list of references. *Emerald 1997* explains that the statistical algorithms in H. Javitz and A. Valdes, "The NIDES statistical component description and justification," Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA March 1994 ("*NIDES 1994*") provides the foundation for the profile-based anomaly detection in *Emerald 1997*:

"Requirements for an anomaly-detection system that became IDES were documented in [6]. This research led to the development of the NIDES statistical profile-based anomaly-detection subsystem (NIDES/Stats), which employed a

wide range of multivariate statistical measures to profile the behavior of individual users [9].”¹²⁹

263. In addition, *Emerald 1997* also directs one to the Network Security Monitor (NSM) system for analysis of network traffic, and identifies L.T. Heberlein, B. Mukherjee, K.N. Levitt., “A Method to Detect Intrusive Activity in a Networked Environment,” Proc. 14th National Computer Security Conference, pp. 362-371, Oct. 1991 (“*Intrusive Activity 1991*”) as describing NSM:

“[T]he Network Security Monitor [7] seeks to analyze packet data rather than conventional audit trails...”¹³⁰

264. Given these explicit references in *Emerald 1997* to both *NIDES 1994* and *Intrusive Activity 1991*, these three references should be considered to be a single disclosure for purposes of anticipation. In the alternative, the citations above provide a motivation to combine these three references in order to make and improve the statistical profiles and network traffic monitoring claimed in the patents-in-suit. As shown in Exhibit U, *Emerald 1997*, *NIDES 1994* and *Intrusive Activity 1991* together satisfied every limitation of the asserted claims of the patents-in-suit (except claim 7, which is invalid as obvious) and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in the patents. Consequently, the noted claims of the patents were not novel when filed.

7. Network NIDES

265. This section covers NIDES and future design work (collectively called “Network NIDES” in this section) as described in the report titled “Next-generation Intrusion Detection Expert System (NIDES) A Summary,” dated May of 1995 (“*Network NIDES*”). SRI International performed this work, with Debra Anderson, Thane Frivold, and Alfonso Valdes (one of the named inventors) listed as the authors of this report. The

¹²⁹ See *Emerald 1997* at 359.

¹³⁰ See *Emerald 1997* at 364.

primary sponsor for this work was the Department of the Navy, Space and Naval Warfare Systems Command (SPAWAR).

266. NIDES was under development from about 1992 through late 1994, with this report delivered in 1995, although NIDES itself was essentially an extension to SRI's Intrusion Detection Expert System (IDES), which dates back to the mid-late 1980s. NIDES was designed to detect intrusive behavior using a combination of anomaly and misuse detection. NIDES was essentially a redesign of the SRI's earlier IDDES, with a focus on a modular and more flexible approach. This modularization and flexibility also helped retarget NIDES to other platforms than its original Solaris audit-based analysis. For example, NIDES was modified to support (1) monitoring of Trusted Xenix systems (Advisor Project), (2) database transaction logs (FBI FOIMS-IDDES Project), (3) and analyzing applications (Safeguard Project).¹³¹

267. The basic NIDES architecture included an anomaly engine performing statistical profiling, a signature/misuse engine, and a resolver. The anomaly detection engine detected previously unknown intrusive activity based on the assumption that this behavior would appear different than normal behavior. The signature/misuse detection engine detected known intrusive activity. The resolver received the analysis from the anomaly and misuse detection engines, integrated multiple reports associated with the same underlying event, and responded by sending a message to the user interface or sending out email.¹³²

268. Network NIDES had the following salient features:

- Monitored network components such as routers and gateways.

¹³¹ *Network NIDES* at 3.

¹³² *Network NIDES* at 22-23.

- Analyzed network packets, audit logs, application logs, and results of other intrusion detection sensors; these data sources were referred to as an event stream.
- Used anomaly detection based on the previously published NIDES algorithms to analyze the event stream.
- Used signature detection to analyze the event stream for known examples of misuse.
- Provided an Application Programming Interface (API) to allow third-party modules and sensors to participate in the system.
- Integrated analysis from an anomaly detection engine, a signature detection engine.
- Supported hierarchical monitoring with domain NIDES monitors reporting to higher-level NIDES monitors.

269. The input stream processed by NIDES was somewhat generalized with 10 general audit record formats supported, another 90 potential types reserved, and 50 allowable user-defined audit types. The primary audit types that SRI worked with were SunOS BSM version 1, SunOS C2, and UNIX accounting records. SRI provided a small application (agen) to allow users to convert their own data sources into a format for NIDES.

270. Figure 10 shows a composite view of the NIDES system as described in *Network NIDES*. This architecture is based on what was described as implemented as of the time of this report as well as what was planned as described in Section 3 “Future Directions.”

271. In the lower left portion of Figure 10 is a description of the basic NIDES monitor responsible for a particular domain (1). It would read an input stream such as network packets (2), and then analyzed that input stream with an anomaly detection engine (3) and a signature detection engine (4). The results of these analyses engines were passed to the Resolver (5), which could integrate the results, take some type of response such as alerting a user (6), or pass its analysis to hierarchically higher monitors

(7). The boxes to the right of the Domain NIDES monitor (1) are simply other Domain NIDES monitors.

272. Above the local domain NIDES monitors is a “higher-level NIDES” monitor (8) that would process the results from the lower level Domain NIDES monitors (1). This higher-level NIDES monitor would be “responsible for the network that supports all the local domains.”¹³³

273. New rules would be developed specifically to detect “suspicious network activity.”¹³⁴ Also, the statistical component would be “easily adaptable to network intrusion detection”¹³⁵ by adapting the subjects (e.g., hosts instead of users) and measures for those subjects (e.g., “the amount of traffic originating or being received” by the subjects).¹³⁶

274. NIDES processed any type of input stream as long as it is mapped to the NIDES audit format via an application called “agen.” In addition to those data sources that were implemented for NIDES, the report also suggested input streams could consist of network packets, reports from other NIDES monitors, “a series of connection attempts,” “the amount of traffic,” reports from the Network Intruder Detector (NID), reports from TCP wrappers about connection attempts, reports from firewalls, SNMP, and reports from application gateways.¹³⁷

275. Furthermore, *Network NIDES* monitors would be considered monitors under any of the parties’ claim constructions.

¹³³ *Network NIDES* at 31.

¹³⁴ *Network NIDES* at 31.

¹³⁵ *Network NIDES* at 32.

¹³⁶ *Network NIDES* at 31.

¹³⁷ *Network NIDES* at 17, 32, 33.

276. In my opinion, as more fully reflected in Exhibit K to my report, the system disclosed in *Network NIDES* satisfied every limitation of the indicated claims of the patents-in-suit and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in the patents. Consequently, the noted claims of the patents were not novel when filed.

277. In addition, as reflect in Exhibit K, the remaining asserted claims of the patents-in-suit are invalid due to obviousness.

278. In particular, it would have been obvious to combine the *Network NIDES* system with other hierarchical intrusion detection systems such as DIDS and GrIDS, and to use their correlation capabilities in *Network NIDES*. It also would have been obvious to combine *Network NIDES* with other systems analyzing network packets and different categories of network traffic data, and to perform NIDES statistical profiling on these categories.

279. In addition, to the extent there were any differences between the configuration of *Network NIDES* and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of *Network NIDES* based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

8. Frank Feather Thesis

280. This section covers Frank Feather's dissertation as described in his thesis "Fault Detection in an Ethernet Network via Anomaly Detectors," dated May of 1992 ("*Feather Thesis*"). This dissertation was based on Frank Feather's research at Carnegie Mellon University and extended the Performance and Anomaly Monitoring System (PAMS) developed primarily by Pamela Hayes. Hayes' work on PAMS was published

as a Senior Thesis at Chatham College, Pittsburgh, PA in May of 1987 ("Characterizing Numeric Time Series Data: Applications to Ethernet LAN Data"). The *Feather Thesis* covered work up through May of 1992. The original PAMS work was published in 1987.

281. The *Feather Thesis* describes a network monitoring system that used an algorithm that was very similar to the Haystack statistical intrusion detection algorithm. The work's primary focus was to quickly detect a wide range of network faults, including those generated by a number of malicious or intrusive activities. The system performed fault detection by applying anomaly detection to a number of measured features of the network traffic. An anomaly was determined by comparing current measurements of network activity to historical measurements for the same measurements; the historical measurements were regularly updated using an exponential aging factor like that used in NIDES. Once anomalous activity was detected, the system further refined the analysis by applying a fault feature vector to the anomaly vector in order to determine the likely underlying cause of the detected anomalous activity.

282. Feather's dissertation had the following salient features:

- Monitored network packets.
- Designed to detect faults in the network ranging from a wide range of causes, including network worms, inappropriate use of the network, mischievous users, or malicious network misuse (e.g., connecting to the network without permission).
- Extended an earlier anomaly detection system called "Performance and Anomaly Monitoring Systems" (PAMS).
- Added the Address Anomaly Detection System (AADS) to extend PAMS' network-wide analysis to host and host-to-host analysis.
- Applied profiles of past activity to determine if currently observed network activity was anomalous.
- Updated profiles using the same exponential decay approach used by SRI's NIDES algorithm.

- Applied a fault feature vector to help determine the cause of the detected anomalous activity.
- Discussed a wide range of network features to measure.
- Deployed and evaluated on an Ethernet network analyzing packets including TCP/IP packets.
- Closely resembled the Haystack intrusion detection system algorithm.
- Suggested also applying the analysis to routers, bridges, file servers, and other network components.

283. The *Feather Thesis* addressed the problem of quick detection and diagnosis of soft failures in a network before they became a noticeable problem to the users. “Soft failures” are generally characterized by degraded performance of the network. In computer security lingo, this is generally referred to as denial of service. “Hard failures” are a complete failure of a component of the network so that no packets can be delivered. Unaddressed soft failures may eventually lead to hard failures. For example, a failing component in the network may lead to intermediate packet loss (soft failure), but that failing component may eventually completely fail leading to hard failure.

284. There are many potential reasons for soft failures in the network. Some of the examples of causes for soft failure listed in the dissertation include inappropriate use of the network, temporary congestion causing delayed transmission, failed host hardware, failure of higher level protocols, mischievous users, Internet worms, or a user connecting a host to a network without permission.

285. This system described in the *Feather Thesis* measured a wide range features about the network as a whole, individual hosts, and pair-wise host communications. Examples of such measurements include packets per minute, load, broadcast packets, packet size, packet type distribution, unexpected hosts on the network, unexpected communication between pairs of hosts, and a wide range of error events.

Such error events include collisions, checksum errors, undersized packets, oversized packets, and dropped packets.

286. Once measurements for the subjects were identified, the system developed expectations of the values for these numbers; that is, it determined the normal values for each of the measurements. The normal values were determined by analyzing past behaviors for the measurements, and the expected normal values were updated regularly using an exponential smoothing of past activities such that recently measured activity were given more weight than past activities – exponential aging.

287. Once profiles of normal values were established, the system compared current measurements to historically normal values for each measurement and identified anomalous measurements.

288. To go beyond simply identifying anomalous activity, the system applied fault feature vectors to the set of anomalous features to automatically identify the likely underlying cause of the observed anomalous behavior. For example, a bridge going down would exhibit one set of anomalous features, while a host using too much memory and paging out memory to a remote disk would exhibit a different set of anomalous features. Fault feature vectors tuned to different fault types could help disambiguate the cause of the detected anomalies.

289. Figure 11 shows the basic algorithm described in this dissertation. The box on the left represents the analysis performed as part of Frank Feather's dissertation (1). At the bottom is the data source, packets, which are fed into the system (2). From the packets a description vector was computed to represent the currently observed network activity. This description vector was composed of two parts: a series of features representing system-wide performance (X_1, X_2, \dots, X_m) (3), and a set of features for individually observed hosts or communicating pairs of hosts ($H_1 \dots H_k$) (4). Together,

these measurements represented the description vector, which is a short-term profile of recent activity.

290. Next, code compared (5) the current description of the network traffic with the profiles (6) of past network traffic. The result of this comparison was an anomaly vector (A_1, A_2, \dots, A_n) (7) describing which of the measured features were anomalous. Once anomalous activity was detected, the system tried to determine the underlying cause of the anomalous activity by applying (8) a weighted fault vector (9) for different fault types to the anomaly vector. The result of this analysis is the “matching score” (10).

291. Interestingly, the approach described in Frank Feather’s dissertation is strikingly similar to Haystack Labs’ Haystack Intrusion Detection System (11).¹³⁸ In Haystack, the primary data source was a system’s audit records (12), and a feature vector for each user was computed (13). An anomaly detection function (14) compared the current feature vector to a profile of historical behavior (15) to determine if there were any anomalous features. The result was an anomaly vector (16) that Haystack called the bernoulli vector. To further determine what type of malicious activity might be causing the detected anomalous activity, Haystack applied (17) a weighted intrusion vector (18) for different classes of intrusive activity, and the result was a weighted intrusion score (19) for each class of intrusion. Haystack took one additional step (not shown) to determine the probability of a given session having a given weighted intrusion score or higher.

292. Under Symantec’s alternative claim construction, the *Feather Thesis* would satisfy the requirement that the long-term statistical profile was an exponentially

¹³⁸ I am familiar with the Haystack statistical algorithm because I analyzed Haystack, including the source code, in order to perform the analysis in my 1994 paper: B. Mukherjee, L.T. Heberlein, K.N. Levitt, “Network Intrusion Detection,” IEEE Network, Vol. 8 No. 3, pp. 26-41, May/June 1994.

aged probability distribution of historical values. However, Feather's short-term statistical description was not aged.

293. The *Feather Thesis* used network packets as its primary data source. The Address Anomaly Detection System (AADS) element of the dissertation, which analyzed hosts and communicating pairs of hosts, processed data coming from the commercial Ethernet-Adapter-Karte (EAK) Monitoring System (EMS) – a network monitoring card that captured the packets and collected the primary statistics.

294. Most individual measurements were analyzed using the Performance and Anomaly Monitoring System (PAMS), which was focused on analyzing univariate time series data. For any given measurement at any given time during the day, a mean and standard deviation were computed. Newly observed data measurements that exceeded the mean by some amount (e.g., more than three standard deviations above the mean) were flagged as anomalous. Furthermore, the AADS component also reported the sudden appearance of a new network address or the sudden drop out of a usually talkative network address.

295. Once anomalous activity was detected, a fault feature vector was applied to the anomaly vector in an attempt to disambiguate the underlying cause of the detected anomalous activity. The dissertation also mentioned the need to automatically correct detected faults, but it did not go into detail as to how that might be achieved.

296. In my opinion, as more fully reflected in Exhibit L to my report, the system disclosed in *Feather Thesis* satisfied every limitation of the indicated claims of the '338 patent, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in the '338 patent. Consequently, the noted claims of the '338 patent were not novel when filed.

297. In addition, as reflect in Exhibit L, the remaining asserted claims of the '338 patent are invalid due to obviousness.

298. The *Feather Thesis* related to the field of network monitoring, which was undergoing significant expansion in the 1992 timeframe in which the thesis was written. By spring of 1993, a standard had been developed for doing hierarchical monitoring. Furthermore, the *Feather Thesis* explicitly directed the reader to use SNMP: "[w]ith SNMP now available on many platforms, it would be very easy to do anomaly detection on data from other network components."¹³⁹ It would have been obvious to combine the work of the *Feather Thesis* with hierarchical monitoring systems, because the entire network monitoring field was already moving in this direction. In addition, the network monitoring SNMP standards for the purpose of network management documented a wide range of network traffic data categories to monitor. It would have been obvious to combine the *Feather Thesis* with these additional known categories of network traffic in order to improve and expand upon its network monitoring capabilities.

299. In addition, to the extent there were any differences between the configuration of the system described in the *Feather Thesis* and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of the *Feather Thesis* based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

9. HP Open View and RFCs

300. This section covers the field of network management and in particular the HP OpenView network management platform and the Internet Engineering Task Force (IETF) standards such as SNMP, MIB, and the Remote Network Monitoring (RMON)

¹³⁹ *Feather Thesis* at 172.

specifications (RMON I and II) (collectively the “RFCs”)¹⁴⁰. The RFCs define an infrastructure and monitoring capability that can monitor a wide range of measurements for many different network elements, determine if those measurements are unusual, and if so, send alerts to higher-level monitors. The SNMP architecture supported a hierarchical design. Furthermore, network management platforms such as HP OpenView provided a wide range of responses for unexpected behavior, including alerting a user through graphical means, audible means, and sending messages to their pagers.

301. Network management applications such as HP OpenView and the RFCs were not classic intrusion detection systems, but they certainly performed “network monitoring” as claimed in the patents-in-suit. Furthermore, the computer security field is simply a part of the overall network management field.¹⁴¹ Many network faults and problems create suspicious behavior patterns that are important for intrusion detection systems to be aware of. Many different intrusion detection systems explicitly relied upon network management systems for part of their functionality. For example, NetRanger used HP OpenView as its user interface. NetStalker could send messages to other network management stations via an SNMP trap message. Both the *JiNao Report*¹⁴² and *Emerald 1997*¹⁴³ explicitly mentioned using SNMP traffic as a data source for their intrusion detection monitors. In addition, all of the MIB and RMON metrics are designed

¹⁴⁰ RFC 1157, A Simple Network Management Protocol (SNMP), May 1990 [SYM_P_0501113- SYM_P_0501142] and [SYM_P_0527111]; RFC 1155, Structure and Identification of Management Information for TCP/IP-based Internets, May 1990 [SYM_P_0501012- SYM_P_0501031]; RFC 1213, Management Information Base for Network Management of TCP/IP-based internets: MIB-II, March 1991 [SYM_P_0501143- SYM_P_0501205]; RFC 1271, Remote Network Monitoring Management Information Base, November 1991 [SYM_P_0501206- SYM_P_0501271]; RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIV2, January 1997 [SYM_P_0603708- SYM_P_0603837].

¹⁴¹ W. Stallings, *SNMP, SNMPV2, SNMPV3 AND RMON I AND 2*, 3rd ed. 1999.

¹⁴² *JiNao Report* at 5-6.

¹⁴³ *Emerald 1997* at 356.

for monitoring network entities such as routers and gateways. Many network attacks such as sweeps and SYN floods are detectable through changes in the standard MIB variables. In fact, almost every network traffic category represented in the patents was represented in the MIBs. Thus, network management applications such as HP OpenView and the RFCs that network management relied upon are directly relevant to the alleged inventions.

302. The network management effort was directed by the Internet Engineering Task Force (IETF), and was worked on by many different organizations and people over the years. The work began in the 1980s, and the first network management documents were published in 1988 (RFC 1067 – SNMPv1 protocol, RFC 1065 – SMIV1 language description, and RFC 1066 – MIB I). Hewlett-Packard developed the HP OpenView network management platform. The HP OpenView for Windows User Guide for Transcend Management Software Version 6.1, which was used for this report, was released in October of 1997 (“*HP OpenView manual*”). However, many other publications document the features of HP OpenView in the relevant timeframe.¹⁴⁴

303. Network management in general and HP OpenView in particular was designed to help “manage large communications infrastructure.”¹⁴⁵ Management included monitoring elements in the network (hosts, routers, bridges, network segments, etc.) for proper or expected behavior, alerting the appropriate entity (a person or another program) when improper or unexpected behavior was detected, and issuing commands to those network elements.

304. The network management paradigm is centered on the concept of a “managed node”, a “management application,” and a standardized interface. The

¹⁴⁴ See, e.g., M. Miller, *MANAGING INTERNETWORKS WITH SNMP*, 2nd Ed., 1997.

¹⁴⁵ M. Rose, *The Simple Book, An Introduction to Internet Management*, 2nd Ed., 1994.

management application monitors, and if necessary, sends controlling information to the managed node. Because the protocols (SNMP) and management information base (MIB) description language (SMI) are standardized, a wide range of management applications and managed nodes can be mixed and matched. Furthermore, a management application can in turn be a managed node for a higher-level management application.

305. Salient features of network management systems prior to November 1997 included:

- Developed by a standards body (Internet Engineering Task Force (IETF)).
- The basic network management architecture provided a two-tiered hierarchy of management application and managed nodes.
- A single manager could monitor many managed nodes, and a single managed node could be managed by many network managers.
- The simple two-tiered network management hierarchy could easily be extended to multiple levels of hierarchy, and this was endorsed with the Manager-to-Manager MIB released in 1993.
- The IETF community endorsed a standard network monitoring capability with the RMON and RMON-II specifications.
- The RMON monitors could be run on stand-alone sensors, routers, or on network management platforms.
- RMON MIBs specified a wide range of subjects to monitor, including network segments as a whole, hosts, host-to-host communications, and communication for a particular application (applications are identified by information in the packet).
- RMON MIBs specified a wide range of network metrics to measure for each subject, including number of packets, number of bytes, and many error conditions.
- RMON sensors could be configured to monitor specific metrics and trigger alerts that were sent to higher-level monitors when unusual conditions are detected.
- HP OpenView provided a published API for third-party applications to interact with the OpenView network management application suite.

- HP OpenView provided a number of additional responses including visual alerts, audible alerts, starting additional programs, and sending a message to a pager.

306. Figure 12, which was scanned in from the 1997 book “Understanding SNMP MIBs” by Perkins and McGinnis,¹⁴⁶ shows the basic approach of network management systems. The basic “management station” and “managed node” can be scaled up to support very large networks. At Level 1 is the basic managed node (1). It receives messages from and transmits messages to the Level 2 manager, which in this case is a dual-role entity (2). As described in the text, the “Dual-role entities are typically used to forward SNMP messages (an SNMP proxy), or to consolidate and synthesize information from many systems and make that information available to a higher-level manager” (3). From our security perspective, the “consolidate and synthesize” means to correlate and integrate. The dual-entity object (2) can then forward the “consolidated and synthesized” information to a 3rd tier management station (4), which can also receive information directly from other managed nodes (5).

307. Figure 13 was scanned in from the 1993 book “SNMP, SNMPv2, and CMIP: The Practical Guide to Network-Management Standards” by William Stallings¹⁴⁷ and was originally part of a 1992 article. Of particular interest for this report is remote monitoring, of which the management is specified by the RMON and RMON-II MIBs. Figure 13 shows an example of an RMON deployment for remote monitoring. In the lower left is a network router supporting an RMON monitor (1). At the bottom middle (2) and the bottom right (3) are PCs running RMON monitors. Each of these monitors must report to a higher-level network management station like the one at the top of the

¹⁴⁶ D. Perkins and E. McGinnis, UNDERSTANDING SNMP MIBS, 1997. (Level 1, 2, and 3 labels and the circled numbers have been added for this discussion).

¹⁴⁷ W. Stallings, SNMP, SNMPv2 and CMIP, The Practical Guide to Network-Management Standards, 1993.

figure (5). The top-level network management station is also running its own RMON sensor (5). Finally, on the far left is yet another RMON monitor (4), but this one not only reports to (5) but it also supports its own local management console.

308. Figure 14 captures several relevant aspects of RMON sensors. In the upper left corner is a sample of several of the subjects an RMON sensor monitored (1). Furthermore, the RMON MIBs describe a wide range of statistics that may be monitored for the various subjects, including the number of packets, bytes, broadcast packets, error messages, and protocol distributions (2).

309. Figure 14 also briefly shows RMON's approach to anomaly detection and alerting (3). The graph's vertical axis represents the number of packets for some subject (e.g., a host) (4). The RMON MIBs (and indeed, the general purpose MIB-II) is highly configurable as to what metrics are monitored for what subjects. The graph's horizontal axis represents time (5). For this measurement and subject, a normal range of values is determined (6), and an upper threshold (7) and a lower threshold (8) bound that normal range. The RMON sensor then monitors the measurement over time, and as long as the measurements remain in the normal range nothing happens. However, when the measurement first exceeds one of the thresholds of normal values (10), the RMON monitor automatically triggers an alert (11) that is sent to a higher-level monitor.

310. For the RMON sensors, the data source is network packets. For example, it can passively monitor all the packets on a network segment. Most network management agents can monitor additional details about the objects they monitor; for example, the standard MIB specification includes details about the state of a particular connection (i.e., a management application can check if a connection is in a half-open connection). Shown in Exhibit X is a chart of the different types of information that could be monitored via SNMP or RMON sensors.

311. Management applications receive messages from managed nodes or from other management stations. In the case of HP OpenView, individual management applications can be started by and receive messages from OpenView.

312. Any individual variable on a managed device can be polled to determine its current value, and the result of that value can be passed to management applications (more then likely, the management application requested that the value be polled). If the managed node observes a specific value for certain variables, it automatically alerts a network management station.

313. Beginning with the Manager-to-Manager Management Information Base (RFC 1451), the normal range of values for a variable can be defined and if the variable exceeds the normal range an alert is automatically sent to hierarchically higher monitors. HP OpenView extends this by providing support for detecting anomalous behavior for categorical variables too.

314. The network management infrastructure provides a wide range of responses to anomalous or suspicious behavior. At the simplest level, simple events (e.g., a network segment failing) can trigger an alert (via an SNMP trap) to be sent to higher-level monitors. Network management objects can also respond to more sophisticated behaviors by setting alarms to be triggered when a value exceeds a normal range.

315. Furthermore, network management applications are expected to be full-blown applications that can provide complex analysis and response. For example, HP OpenView provided a graphical user interface (GUI) that could automatically alert a user, sound an audible alarm, automatically execute additional programs, or even activate a remote paging device. Finally, network management systems that serve in a dual-role capacity could alert a hierarchically higher network management system. For example,

Perkins says dual-role entities are used to “consolidate and synthesize information from many systems and make that information available to a higher-level manager.”¹⁴⁸

316. As noted in Exhibit M, the *HP OpenView manual* specifically referenced and relied upon the noted RFCs. Furthermore, network management applications such as HP OpenView were publicly used with all of the noted RFCs. In the alternative, given the explicit direction in the *HP OpenView manual* to use the RFCs for network management, such a combination would have been obvious.

317. In my opinion, as more fully reflected in Exhibit M to my Report, the system disclosed in *HP OpenView manual* satisfied every limitation of the indicated claims of the ‘203 and ‘615 patents, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the asserted claims of the ‘203 and ‘615 patents (except ‘615 claims 7 and 64-73) were not novel when filed.

318. Furthermore, network management applications such as HP OpenView were in public use with the RFCs for network management prior to November 9, 1997, and this public use also anticipates the indicated claims.

319. In addition, as reflect in Exhibit M, it would have been obvious to combine the statistical detection method from the *Feather Thesis* with the HP OpenView system and Internet Standards, for the reasons discussed previously regarding the *Feather Thesis*, and thus claim 7 of the ‘615 patent is invalid due to obviousness. Given the fact that HP OpenView managed network infrastructure, and a firewall is network infrastructure, it would have been obvious to treat a firewall as a managed node. Thus, ‘615 claim 64-73 are invalid due to obviousness.

¹⁴⁸ Understanding SNP MIPs, p. 7.

320. In addition, to the extent there were any differences between the configuration of HP OpenView system and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of HP OpenView based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

10. NetStalker

321. This section covers the commercial product NetStalker sold by Haystack Labs, Inc., as described in the May 1996 document "NetStalker Installation and User's Guide: Version 1.0.2" (*"NetStalker manual"*). Haystack Labs, Inc. developed NetStalker, and the primary developers were Steve Smaha, Steve Snapp, Jessica Winslow, Richard Letsinger, Crosby Marks, Charisse Castagnoli, Brita Womack, and Kristin Johnson.

322. NetStalker enhanced the Network Systems Corp. (NSC) routers by providing richer signature analysis to look for a wider range of misuse than is possible through the NSC router rules. NetStalker also could alert users through a number of means and could automatically modify router filter rules to block detected attacks. Examples in the manual of misuse detected by NetStalker included illegal FTP sessions, SATAN scans, IP spoofing, ICMP error messages, and fragmented TCP headers.

323. NetStalker was a software package that the user installed on a workstation. The NetStalker software was configured to monitor data from and control one or more NSC routers. The NSC routers needed to be configured to send messages to the NetStalker host and accept configuration changes from the NetStalker host. The configuration changes to the NSC routers might include selecting a different set of packets to forward to the NetStalker host or blocking a particular address. NetStalker could also be configured to send messages to other network management stations via an

SNMP trap message. Because the SNMP standards already supported hierarchical architectures, NetStalker systems could be part of a large-scale, hierarchical security management system.

324. Figure 15 shows the basic configuration. At the bottom is an NSC router (which can also operate in bridge mode) (1). The NSC router observed packets, applied a set of filters to each packet, and, depending on the result of those filters, sent a selection of those packets to the NetStalker host (2).

325. The NetStalker system applied a series of filters, where each filter implemented Haystack's pattern matching signature analysis engine (3). If an event passed through all the filters, it reaches the alarm element of NetStalker (4). The alarm element applied simple correlation and a threshold to determine when to take an action such as alerting a user (5) or sending a trap message (6) to a hierarchically higher network management system (7).

326. The original source of data for NetStalker was packets. The NSC router would read the packets (acting as a router or a bridge), select some subgroups of those packets, and forward those packets to the NetStalker host. From those packets, NetStalker could extract numerous metrics from the packets including protocol types (e.g., IP, ICMP, TCP), services (e.g., Telnet, FTP), a wide range of ICMP error messages (e.g., destination unreachable, source quench), and network addresses.

327. NetStalker provided a wide range of responses including: popping up an alarm window on the GUI, sending an SNMP trap message to a network management station, emailing a user, modifying the NSC routers to block an attacker's IP address, and dialing a user's pager.

328. As noted in Exhibit N, NetStalker generated snmp traps, and HP OpenView was designed to use such SNMP traps. Furthermore, both systems were on-sale prior to November 9, 1997. Therefore, the use of the two together constitutes a public use. In the alternative, given that the NetStalker product generated snmp traps and

the *NetStalker manual* stated “[f]or each alarm generated by NetStalker, you can configure one or more alarm handlers to serve as communications channels from NetStalker to you, to other network management tools, or to respond to the alarm”¹⁴⁹ such a combination would have been obvious.

329. In my opinion, as more fully reflected in Exhibit N to my report, the NetStalker product in use with HP OpenView satisfied every limitation of the indicated claims of the ‘203 and ‘615 patents, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the claims of the ‘203 and ‘615 patents were not novel when filed.

330. In addition, as reflect in Exhibit N, it would have been obvious to combine the NetStalker system with the HP OpenView system and RFCs, for the reasons discussed previously, and thus the indicated claims are invalid as obvious. ‘615 claim 7 is similarly obvious for the reasons discussed previously with regard to HP OpenView.

331. In addition, to the extent there were any differences between the configuration of NetStalker and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of NetStalker based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

11. NetRanger

332. I have reviewed in detail the NetRanger documentation listed on Exhibit O. In addition, I spoke to Mr. Daniel Teal in order to confirm my understanding of the functionality of NetRanger. I also reviewed Mr. Teal’s expert report and agree with his opinions. I adopt as part of my own report Mr. Teal’s description of the NetRanger

¹⁴⁹ *NetStalker manual* at 4-2.

product as well as the rest of the analysis in his expert report, including the Figures and the NetRanger invalidity chart.

333. Under Symantec's alternative claim construction of "monitor" the NetRanger system is still anticipatory. The NetRanger NSX was software that could be reconfigured, for example, to add additional signatures. The NetRanger Director could be reconfigured to consolidate multiple alarms into a single alarm based upon user preferences. The NSXs collected, analyzed and responded to suspicious network activity – using analysis engines and functionality to respond to the detection of events by sending alarms. The NetRanger Director's SAP (Security Analysis Package) correlated events based upon its SQL queries. The NetRanger Director also provided response by alerting the user about problems.

334. In my opinion, as more fully reflected in Exhibit O to my report, the NetRanger product and the NetRanger User's Guide Version 1.3.1 satisfied every limitation of the indicated claims of the '203 and '615 patents (except '615 claim 7), and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims of the '203 and '615 patents were not novel when filed.

335. I also agree with Mr. Teal's opinion that it would have been obvious to combine NetRanger with a statistical intrusion detection system if one's primary purpose was a research system. Thus, as reflected in Exhibit O, '615 claim 7 and the '212 claims are invalid as obvious.

336. In addition, to the extent there were any differences between the configuration of NetRanger and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of NetRanger based upon the nature of the problem to be solved. Such configuration changes would have been

motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

12. ISS RealSecure

337. RealSecure, as described in the various documents listed in Exhibit P, was a network-based intrusion detection and audit system. It used a signature-based engine to detect a wide range of attacks (e.g., SYN floods), and could generate detailed network audit logs of activities such as which web pages were requested, what users logged into FTP servers, and what files they transferred. The manual recommended placing RealSecure at the organization's gateway to the Internet, either behind the organization's firewall or Internet router. A single RealSecure management console could configure and receive reports from multiple RealSecure sensors spread across the organization.

338. In my opinion, as more fully reflected in Exhibit P to my report, the RealSecure product satisfied every limitation of the indicated claims of the '203 and '615 patents, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims of the '203 and '615 patents were not novel when filed.

339. In addition, as reflect in Exhibit P, the remaining asserted claims of the '203 and '615 patents are invalid due to obviousness. To the extent there were any differences between the configuration of RealSecure and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of RealSecure based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

13. Network Level Intrusion Detection

340. This section covers the University of New Mexico's Department of Computer Science effort to develop a network-based intrusion detection system. The work is described in an August 1990 paper published under the title "The Architecture of a Network Level Intrusion Detection System" ("*NLID 1990*"). The authors were primarily interested in detecting and reacting to computer worms. While aware of misuse detection techniques, the authors wanted to focus on anomaly detection. The metrics their system collected included timestamps (which also provided time of day information), packet length, packet drops, source-destination pairs, and protocol type. These basic metrics could be combined into a wide range of other metrics (e.g., histograms of packet size distributions), and the authors were concerned with determining which metrics were good for intrusion detection and which metrics were not. To perform anomaly detection the authors used a generic algorithm classifier. The classifier would develop a rule base that represented normal behavior so that anomalous behavior could be flagged. The rule base of historical behavior would be periodically updated.

341. In my opinion, as more fully reflected in Exhibit Q to my report, *NLID 1990* satisfied every limitation of the indicated claims of the '338 patent, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims were not novel when filed.

342. In addition, as indicated in Exhibit Q, the remaining indicated claims of the '338 patent are invalid as obvious. To the extent there were any differences between the configuration of *NLID 1990* and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of *NLID 1990* based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

14. '750 Thompson patent

343. This section describes U.S. Pat. No. 5,825,750 "Method and Apparatus for Maintaining Security in a Packetized Data Communications Network," filed by Horace Thompson on March 29, 1996 (the '750 patent). This patent covers packet-based networks; the embodiment of the invention described in the paper is an ATM-based network. However, one of ordinary skill in the art would have understood from the disclosure that the methods described could also be applied to a TCP/IP network given that TCP/IP commonly runs over ATM networks. The '750 patent includes applying both statistical and rule-based analysis to network packets in order to detect possible security problems. For the anomaly analysis, profiles were generated and periodically updated. Packet metrics that were measured included bandwidth measurements and error conditions. Once a potential security problem had been detected, the sensor intrusion detection system could impose itself in the middle of the communication (man in the middle) to protect the secure node, request additional information, or prevent the potential attacker from making a connection with a target system.

344. In my opinion, as more fully reflected in Exhibit R to my report, the '750 patent satisfied every limitation of the indicated claims of the '338 patent, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims were not novel when filed.

345. In addition, as indicated in Exhibit R, the remaining claims of the '338 patent are invalid as obvious. To the extent there were any differences between the configuration of the '750 patent and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of the '750 patent based upon the nature of the problem to be solved. Such configuration changes would have been

motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

15. Stake Out

346. This section covers the commercial product called "Stake Out" sold by Harris Corporation. This information is taken from their "Stake Out Network Surveillance White Paper." Stake Out was a network-based intrusion detection that ran on a Sun Sparc 20 computer running Solaris 2.5. It would sniff network packets and process TCP/IP networks using a combination of anomaly detection and misuse detection techniques to detect suspicious behavior. Some of the metrics measured by the system included number of packets, type of packets, and source-destination traffic patterns. Once potentially suspicious behavior was detected, the system would respond by alerting the user, modifying the sensor to collect additional information, and potentially send an alert to a hierarchically higher monitor via an SNMP trap message.

347. In my opinion, as more fully reflected in Exhibit S to my report, Stake Out the product and the printed publication white paper satisfied every limitation of the indicated claims of the '338 patent, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims were not novel when filed.

348. In addition, as indicated in Exhibit S, the remaining claims of the '338 patent are invalid as obvious. To the extent there were any differences between the configuration of Stake Out and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of Stake Out based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

16. Automated Alarm System

349. This section covers Los Alamos National Laboratory (LANL) Automated Information System, as described in: W. Huntman, "Automated Information System - (AIS) Alarm System," Proc. of the 20th National Systems Security Conference (Oct. 1997) ("*AIS*"). The LANL Alarm System was primarily a hierarchical architecture to integrate a potentially wide-range of security related sensors. The core architecture was a 2-tier architecture where individual sensors sent reports to a "Central Assessment" center, which correlated the reports and optionally took a number of responses including: directing the sensors to modify their analysis, send an email message, close connections, and notify hierarchically higher Central Assessment systems. One example sensor the paper mentioned could be a network sniffer that reported specific patterns of misuse or anomalous patterns of behavior. A Central Assessment component could act as a "Sensor for the next higher level Alarm System", creating a hierarchical monitor. The Alarm System would also provide a standard API for all the components to communicate.

350. In my opinion, as more fully reflected in Exhibit T to my report, *AIS* satisfied every limitation of the indicated claims of the '203 and '615 patents, and enabled one of ordinary skill in the art prior to November 9, 1997 to practice the alleged inventions claimed in these patents. Consequently, the indicated claims were not novel when filed.

351. In addition, as indicated in Exhibit T, the remaining claims of the '203 and '615 patents are invalid as obvious. To the extent there were any differences between the configuration of *AIS* and the configuration described in the patents-in-suit, it would have been obvious to modify the configuration of *AIS* based upon the nature of the problem to be solved. Such configuration changes would have been motivated by a desire to address the problem of detecting and responding to suspicious network activity in very large enterprise networks.

XI. OBVIOUSNESS

352. In the previous sections, I have already explained many reasons why one of skill in the art would have been motivated to combine a particular reference or system with another. In addition to those explanations, I will further elaborate some factors that relate to the obviousness of the alleged inventions. In general when addressing obviousness, I have considered the issue of what a person of ordinary skill in the art prior to November 9, 1997 would have done if confronted by the problem of detecting network intrusions if that person had no knowledge of the patents-in-suit. I have further considered the problem of detecting network-related intrusions across an enterprise network, spanning multiple domains, and how that problem would have been confronted by one of skill in the art in the relevant timeframe.

353. It is important to point out that by November 1997, there was a very rich body of prior art in the network monitoring and intrusion detection fields. Many different individuals and research groups had been working on these issues for many years. Thus, one of ordinary skill in the art would have known about many such systems and past work. Given this large body of prior work, it is impossible for me to enumerate all the possible, likely combinations. My report discusses merely a sample of the possibilities.

354. As discussed earlier with regard to *Emerald 1997*, much of the motivation to look at particular categories of network traffic stems directly from the types of attacks that can be launched against networks. As I understand the alleged inventions of the patents-in-suit, a key issue is whether one of skill in the art would have been motivated to use the claimed network traffic data categories as an input into a network-based intrusion detection system. In my opinion, the claimed network traffic data categories broadly direct one to look at parts of mainly network packet headers, or to look at the volume of data in a packet. If looking for network-based intrusions, or indeed any type of activity

on a network, one of skill in the art prior to November 1997 would naturally have looked at these parts of a network packet.

355. As I described earlier, a network packet typically consists of a header portion and a data portion. The structure of network packets was well-known and indeed trivial in the computer science field in the 1990s. In addition, monitoring and parsing different portions of network packets to gain relevant information about a particular part of a network packet was similarly common in the computer science field.

356. By 1997, it was well-known that attackers could use either the header portion of a packet or a data portion of a packet to carry out an attack.¹⁵⁰ While attacks like SYN floods, IP scans, port scans, and SATAN scans primarily relied on the header portion of network packets, other attacks such as buffer overflow (also known as buffer-overrun) exploits, password cracking, and worms relied heavily on transporting data or code within the data portion of network packets.

357. Attacks that relied primarily on the header portion of network packets used the various data fields contained in the header to carry out the attack. These fields are ordinarily used by various network protocols to administer the transfer of data via network packets. For example, some header fields—such as the SYN and ACK fields in a TCP header—relate to administering the establishment of network connections. Other header fields—such as the source IP address and destination IP address fields in an IP header—relate to administering the transfer of data from one computer to another over a network. Because well-known attacks used such fields in malicious ways, one of ordinary skill in the art would have been motivated to monitor header fields using both profile-based anomaly detection and signature detection.

¹⁵⁰ NetRanger 1.3.1 User's Guide at 4-61 (see, e.g., IP Fragment attack and FTP CWD ~root exploit).

358. Other attacks relied primarily on the data portion of network packets to transfer malicious data or code. As noted previously, it was well-known by the early 1990s that several network-based attacks transported malicious code, data, or commands via the data portion of network packets. Because these attacks placed malicious code, data, or commands in the data portion of network packets, one of ordinary skill in the art would have been motivated to monitor the data portion of network packets using both profile-based anomaly detection and signature detection.

359. As I described earlier, network protocols are typically organized as a layered stack, with each protocol being built on the protocol or protocols directly beneath it. It was well-known by 1997 that attacks could be carried out at any layer of the network protocol stack. For example, while SYN attacks were carried out using transport layer TCP flags relating to connection establishment, many attacks directed at FTP servers were carried out using application layer FTP-specific commands. Because each and every network protocol layer was the subject of attacks, one of ordinary skill in the art would have been motivated to monitor each and every layer using both statistical profile-based anomaly detection and signature detection.¹⁵¹

360. To summarize, in order to increase the number of attacks that could be detected, one of ordinary skill in the art would have been motivated to use intrusion detection systems to analyze all portions of a network packet, including the header and data portions of packet formats from each network protocol layer. The nature of the problem to be solved would have led one of skill to examine the claimed network traffic data categories.

¹⁵¹ The need to monitor each and every network protocol layer is related to the need to monitor both the header and data portions of network packets, since packets at a given layer typically encapsulate packets at a higher layer. For example, in order to monitor transport-layer TCP segments, one has to extract those segments from the data portion of network-layer IP packets.

361. By 1997, it was also well-known to those of skill in the art that attackers targeted all types of network entities, including specialized entities like routers, gateways, proxy servers, firewalls, proxy servers, and any other device connected to a network. There were several reasons why attackers targeted a wide variety of network entities. First, because of the prevalence of use of standard network protocols, an attack that exploited one of these standard protocols could typically be launched against the wide array of network entities using that protocol. For example, a SYN flood that was initially developed to attack servers utilizing TCP could just as easily be launched against a handheld networked device that also utilized TCP. Thus, just as network-based intrusion detection leveraged the network protocols to monitor several different types of network entities, attackers could also leverage the use of these protocols to expand their base of attack targets.

362. Second, attackers had an incentive to attack these specialized network entities that comprised the internet infrastructure itself. Infrastructure entities like routers and gateways were attractive targets because an attack on those targets would be felt by every network entity that used the router or gateway as an intermediary for communication to the rest of the network. In other words, attacks on routers and gateways potentially would be more widely felt. Because of the importance of these entities that comprised the network infrastructure itself, one of ordinary skill in the art would have been motivated to protect these entities using intrusion detection systems.

363. By November 1997, there also existed a strong motivation to design intrusion detection systems to interoperate with both other intrusion detection systems and other types of network security devices. Standards-based efforts such as CIDF had emerged to promote interoperability through the use of common protocols and common APIs. There were several factors driving this move towards greater interoperability. From a technical standpoint, greater interoperability enhanced the ability of an intrusion

detection system to analyze information from a wider range of network security devices (like firewalls). Greater interoperability also enhanced scalability by allowing intrusion detection systems to work and communicate with each other regardless of whether these systems were administered by the same party. In this way, cooperation enabled by greater interoperability would allow intrusion detection systems to scale up to larger networks. From a commercial standpoint, greater interoperability allowed new products to work with and leverage existing deployments of heterogeneous network security solutions.

364. Both the SRI IDES/NIDES/EMERALD team and UC Davis' Computer Science Laboratory were well-known players in the intrusion detection space in the 1990s. Both groups published, presented publicly at conferences, and were involved in related DARPA projects. One of ordinary skill in the art would have known about these groups, and would have considered them to be doing related work. Given that both teams were attempting to solve problems relating to network intrusion detection, it would have been obvious to one of skill in the art that it would be worthwhile to combine systems from these two groups.

365. In addition, given the number of "surveys" and compilations discussing various different intrusion detection systems in existence prior to November 9, 1997, there was ample direction and motivation to combine such systems. One such example¹⁵² is a 1994 paper on which I am a named author: B. Mukherjee, L. Todd Heberlein and K. N. Levitt, "Network Intrusion Detection," IEEE Network May/June 1994 ("*NID 1994*").¹⁵³ This paper notes that "[t]he intrusion detection problem is becoming a

¹⁵² For additional examples in 1997, see L. Todd Heberlein, Network Radar presentation, 24 July 1997; and

<http://web.archive.org/web/19971011083618/www.hokie.bs1.prc.com/ia/N2-TODD.htm>

¹⁵³ I understand that a copy of this publication was actually produced from the files of Mr. Porras, one of the named inventors. See SRI 058251.

challenging task due to the proliferation of heterogeneous computer networks.”¹⁵⁴ The paper also points out that it is common to combine statistical and rule-based systems, and provides the reader with several examples of each: “[t]ypically, IDSs employ statistical anomaly and rule-based misuse models in order to detect intrusions...” “[i]n this paper, several host-based and network-based IDSs are surveyed, and the characteristics of the corresponding systems are identified.”¹⁵⁵ The paper encourages further investigation: “new and more-effective detection strategies must be investigated.”¹⁵⁶ The paper specifically encourages investigation into intrusion detection for large networks: “much more research is expected to be conducted, e.g., how can the intrusion-detection concept be extended to arbitrarily large networks...”¹⁵⁷

366. Given these specific pointers in *NID 1994* to a variety of existing intrusion detection systems, as well as the specific direction to combine statistical and rule-based approaches and expand IDS to better cover arbitrarily large networks, one of skill in the art would have been motivated to combine the systems discussed in this paper, as well as other related systems, to achieve the goals of the paper. Attached in Exhibit V is a chart demonstrating how the references and explanations of different existing systems in *NID 1994* anticipate many of the asserted claims of the patents-in-suit. In the alternative, the *NID 1994* paper makes it obvious to combine all these systems. As further indicated in Exhibit V, the remaining asserted claims are rendered obvious in combination with the indicated references. It would have been obvious to combine the disclosure in *NID 1994* with these additional systems for the reasons discussed previously. *NID 1994* encouraged one of skill to investigate and combine existing systems.

¹⁵⁴ *NID 1994* at 26.

¹⁵⁵ *NID 1994* at 26.

¹⁵⁶ *NID 1994* at 41.

¹⁵⁷ *NID 1994* at 41.

A. SECONDARY CONSIDERATIONS

367. I have been informed that it is appropriate in analyzing the obviousness of an alleged invention to consider “secondary considerations” of non-obviousness. I understand that secondary considerations include: commercial success of the invention; satisfying a long-felt need; failure of others to find a solution to the problem; and copying of the invention by others. Other secondary considerations include licensing by competitors and contemporaneous recognition of the inventor’s achievements.

368. Based upon my review and understanding of the facts, as well as my own personal knowledge, I am not aware of any secondary considerations supporting a finding that the patents-in-suit were not obvious.

369. In my opinion, the Emerald system disclosed in the patents-in-suit did not satisfy any long-felt need that I was aware of. I was very familiar with intrusion detection systems in the late 1990s and was deeply involved in the field. I do not recall anyone identifying any “need” fulfilled only by the Emerald system. My peers in the intrusion detection field were aware of the Emerald system through conferences, but I do not recall any particularly unique feature of Emerald that was considered extremely valuable to the intrusion detection community.

370. Furthermore, I am not aware of anyone copying the Emerald system.

371. Many other systems, some of which have been discussed in this report, were able to successfully detect computer intrusions and in particular network attacks. For example, the NSM system and its many different incarnations, such as ASIM, were successful in detecting network attacks. (NSM became ASIM, which as indicated in Mr. Teal’s report was widely used by the Air Force). Thus the government, which funded Emerald, was actually using other intrusion detection systems, including some that started out as research-oriented projects such as NSM. In addition, DIDS, another

research-oriented system from UC Davis, was very successful in developing useful correlation features to track users across different systems, solving a key problem in computer security.

372. It is important to recognize that there is a difference between actual commercial products and research funded by the government. Laudatory statements regarding research projects do not necessarily translate into success in the commercial world, because research projects are evaluated on a different set of metrics than actual commercial projects. Research systems typically are run in a limited, lab environment, whereas commercial systems need to run robustly in a messy, real-world environment. Furthermore, research projects typically only address a small part of the problem, whereas commercial systems need to have a rich supporting infrastructure to make them usable products.

373. In the commercial realm, the NetRanger system was a commercial success used by Fortune 500 companies to protect their networks from intrusions. NetRanger was also successful in government testing, with the DOD/SPOCK report stating:

“Results of the tests clearly demonstrated that when properly configured, the NetRanger hardware/software package:

- 1) Can be used to detect, report, and act on intrusion related activities launched across a network with a high degree of accuracy,
- 2) Would detect all attempted penetrations signatures contained in the default list as installed in the NetRanger for this demonstration,
- 3) Can be used to provide practical and effective intrusion detection, reporting, and selected automatic response actions.”¹⁵⁸

¹⁵⁸ DOD/SPOCK Report at 2 [SYM_P_0074255- SYM_P_0074481 at SYM_P_0074263].

In addition, the DOD/SPOCK report concluded “[i]n the true sense, this suite of tests proved the viability of Real-Time Network Intrusion Detection and Response for implementation today, in a warfighter networked environment.”¹⁵⁹

374. As noted in Mr. Teal’s report, the government also actually purchased and used the NetRanger product. I do not believe that the EMERALD system ever achieved the success that NSM/ASIM and NetRanger did.

XII. INVENTORS’ FAILURE TO DISCLOSE BEST MODE

A. Legal standard

375. I understand that the patent laws require that if an inventor knows of a best mode of practicing the claimed invention, the inventor must disclose that best mode. I also understand that the legal standard for this best mode requirement involves two factual inquiries: (1) a subjective determination of whether the inventor had a best mode of practicing the claimed invention; and (2) if the inventor had a best mode of practicing the claimed invention, an objective determination of whether the best mode was disclosed in sufficient detail to allow one skilled in the art to practice it. I have been asked to render an opinion regarding both of these inquiries.

B. Documents relied upon and methodology

376. As part of my inquiry into whether the inventors satisfied the best mode requirement, I have reviewed many documents. For the purposes of this litigation, SRI has placed into escrow a computer containing source code¹⁶⁰ related to the patents-in-suit.

¹⁵⁹ DOD/SPOCK Report at 5.4 [SYM_P_0074255- SYM_P_0074481 at SYM_P_0074287].

¹⁶⁰ In this report, I am using the term “source code” to encompass not only the text files that are compiled or interpreted in order to generate an executable software component, but also any associated initialization files, configuration files, or other types of input files that are read in by the software component.

Although the sheer number of source code files made it impossible for me to analyze every single source code file on the computer, I have performed numerous electronic searches across the entire escrowed computer and also extensively browsed the file system of that computer. Based on this searching and browsing, I identified many relevant files and extensively reviewed those files. I have also reviewed various documents produced by SRI in this litigation, various SRI publications, the patents-in-suit, the *Live Traffic Analysis* and *Statistical Methods* articles incorporated by reference into the patents-in-suit, and the appendix of source code filed with the patents-in-suit.

377. In addition, in order to understand the inventors' subjective beliefs as to their best mode of practicing the alleged inventions at the time of filing, I have reviewed emails from the inventors in the relevant timeframe. I have also reviewed transcripts of the inventors' deposition testimony as well as testimony from others working on the Emerald project at the time (furthermore, I actually attended many of these depositions).

378. In my review of the source files that SRI placed into escrow, I made use of the RCS and CVS software systems that were provided with the escrow computer and by which some of the source code files were maintained.¹⁶¹ Both RCS and CVS are software systems known as "version control" systems. Such systems are also commonly referred to as "revision control" systems. Version control systems are frequently used in environments where (1) it is important to be able to retain previous versions of files; and (2) several different users are making changes to the same set of files. Version control systems are often used by software development teams. For each version of a file, a version control system typically maintains a copy of that version and information associated with that version, including the date that version was made, who made the

¹⁶¹ For a description of CVS, see "Version Management with CVS," available at <http://ximbiot.com/cvs/manual/>. For a description of RCS, see "Official RCS Homepage," <http://www.cs.purdue.edu/homes/trinkle/RCS/>.

version, and any comments that the author of the version wrote to describe that version. Earlier versions of files and the information regarding those earlier versions can be retrieved at a later point in time. Both RCS and CVS provide the above functionalities.

379. For the files maintained by RCS or CVS, I used RCS commands or CVS commands respectively to determine the dates associated with various versions of the files. For the files not maintained by RCS or CVS, I used standard Unix commands to view the last modification date of the files.

C. The eXpert signature engine, SRI's best mode for analyzing network traffic data, was not disclosed

380. To determine whether the inventors had a best mode for practicing the claimed inventions, I first reviewed the source code underlying several SRI software components that were developed prior to November 9, 1998. One such software component for which I reviewed the underlying source code was a signature engine developed by SRI named eXpert (I believe this is pronounced "e-expert").¹⁶² Among the source code files underlying eXpert, there were several files containing P-BEST rules. Because the rules in these files acted as signatures for analyzing network traffic data, I will use the term "P-BEST rule" and "P-BEST signature" interchangeably. Before I

¹⁶² I understand that Mr. Porras has testified that the source code file eXpert.c, source code files from the "libunit" library, source code files from the "libmessage" library, and source code files from the "libapi" library are all part of the set of the source code files that form the code base for eXpert. See Porras 30(b)(6) Depo. Tr. 193:23-195:19. To the extent that these files were included on the escrowed computer or in the appendix to the patents-in-suit, I reviewed those files. I also reviewed several source code files that contain core functionalities of eXpert but that Mr. Porras failed to mention. Mr. Porras has also testified that a "knowledge base" is also part of eXpert. Porras 30(b)(6) Depo. Tr. 195:2-13. Based on my review of the source code underlying eXpert, I have ascertained that the "knowledge base" that Mr. Porras alluded to in his testimony is a suite of P-BEST rules described later in my report.

explain how the various subcomponents of eXpert work, I will explain briefly the use of expert systems in misuse or signature detection engines.

1. Expert systems

381. An expert system is a computer program that uses a set of rules to derive conclusions. I am familiar with several types of expert systems and at trial I may testify regarding these systems. The rules used by expert systems are often called “inference rules.” Inference rules are of the form “if X then Y.” An example would be, “If you are eligible to vote, then you are at least 18 years of age.” Note that these rules can be much more complex and involve several Boolean operators. An example of a slightly more complex rule would be: “If X and Y, then Z.” Some expert systems employ a technique known as “forward chaining.” Forward chaining is the process of applying inference rules to facts in order to reach conclusions or “goals.”

382. Signature engines are frequently implemented using an expert system. Note that the expert system itself does not typically provide any signatures. Domain specific signatures must be written using the expert system’s rule grammar in order for the expert system to be applied to a particular domain.

2. SRI’s P-BEST expert system

383. There are many well-known, publicly available expert systems.¹⁶³ Since the 1980s, SRI has used an internally developed forward-chaining rule-based expert system called P-BEST.¹⁶⁴ Based on my review of SRI publications, it is my understanding that in 1998 SRI viewed P-BEST as superior to other forward-chaining

¹⁶³ Perhaps the most well-known, publicly available expert system is CLIPS. See <http://www.ghg.net/clips/CLIPS.html>.

¹⁶⁴ See “Detecting Computer and Network Misuse Through the Production-Based Expert System Toolset (P-BEST),” Proceedings of the 1999 IEEE Symposium on Security and Privacy, Oakland, California, May 9-12, 1999 [SYM_P_0068713-28] (hereinafter “P-BEST 1999”) at 6-7.

rule-based expert systems, and had consistently used P-BEST in its intrusion detection projects (including IDES, NIDES, and EMERALD).¹⁶⁵ In touting P-BEST in October 1998, Mr. Porras (one of the inventors) and his colleague Mr. Lindqvist pointed to its speed, its easy-to-use grammar for defining rules, and its close integration with the C programming language.¹⁶⁶ P-BEST consisted of both runtime libraries and a translator that translated rules written in the P-BEST rule grammar to C code. This translation allowed the rules to be applied at runtime with increased speed. The speed with which the rules could be applied was an extremely important factor in SRI's choice to use P-BEST to implement eXpert, because the P-BEST signatures would need to be applied to real-time network traffic.

3. SRI had developed a fully-functional version of eXpert prior to November 9, 1998

384. I have analyzed the version of eXpert that SRI had developed as of November 9, 1998. Specifically, I have analyzed the C source code underlying eXpert—including the source code underlying P-BEST—and the suite of P-BEST rules that SRI had developed for eXpert. I have also analyzed the version of estat that SRI had developed as of November 9, 1998 in order to compare eXpert with estat.¹⁶⁷ Based upon my review of the source code, deposition testimony, and communications from the inventors described in this section, it is my opinion that as of November 9, 1998 the inventors considered eXpert/P-BEST to be their best mode of practicing network

¹⁶⁵ See P-BEST 1999 at 3.

¹⁶⁶ P-BEST 1999 at 1. Although this paper was published in 1999, I have reviewed a draft of this paper dated October 1998 that expresses the same praise for P-BEST expressed in the published version. See SRI 151720-40.

¹⁶⁷ As discussed previously in my report, estat is a generic code for implementing statistical profiling. I have called it “generic” because the estat code included in the appendix is not specific to analyzing network traffic, but instead is designed to accept any type of data as long as it is formatted into a structure estat can understand.

surveillance / network monitoring / event monitoring for detecting “suspicious network activity” as claimed in the patents-in-suit.¹⁶⁸ Furthermore, the inventors also considered eXpert/P-Best to be their best mode of practicing “a signature matching detection method” as claimed in the ’212 patent.

385. I understand that various SRI witnesses have given conflicting testimony regarding the eXpert signature engine. There has been conflicting testimony regarding the extent to which code for eXpert was included in the appendix to the patents-in-suit.¹⁶⁹ There has also been conflicting testimony regarding the functional state of eXpert as of November 9, 1998.¹⁷⁰ For example, Mr. Porras has testified that he was not sure when SRI first integrated P-BEST with the rest of eXpert.¹⁷¹ Despite this conflicting testimony, it is clear from my review of relevant documents and source code that as of

¹⁶⁸ Several documents indicate that SRI had implemented versions of eXpert by the July 1998. *See* Email from P. Porras to K. Skinner and U. Lindqvist entitled “emerald-ftp delivery” (July 7, 1998) [SRIE 0400354-55] (referencing compiled version of eXpert); Email from P. Neumann to Cathy Schott entitled “Re: FW: FY99 Incremental Funding (September 3, 1998) [SRIE 0020286-88](stating that SRI “had now completed integration of the EMERALD eXpert as well as the EMERALD statistics component”); SRI Source Code REQUEST # 3.

¹⁶⁹ Porras 30(b)(6) Depo. Tr. 212:25-216:11.

¹⁷⁰ Porras 30(b)(6) Depo. Tr. 212:25-216:11.

¹⁷¹ Porras 30(b)(6) Depo. Tr. 190:5-23, 212:25-216:11. Because Mr. Porras’s testimony alone does not provide an adequate basis for me to determine whether he believed at the time of filing that he possessed a best mode for practicing the claim inventions, I have looked to documents produced by SRI in this litigation that are contemporaneous with or prior to the filing of the shared specification underlying the patents in suit. Specifically, I have reviewed an email sent by Mr. Porras to Mr. Lindqvist on November 9, 1998, the exact date of the filing of the application underlying the patents-in-suit. *See* Email from P. Porras to U. Lindqvist entitled “Re: 1999 IEEE Symposium on Security and Privacy Paper Submission (fwd)” (November 9, 1998) [SRIE 0275123]. In this email, Mr. Porras informs Mr. Lindqvist about the work that SRI had performed over the prior two weeks in analyzing test data provided by MIT’s Lincoln Labs as part of the DARPA 1998 offline-evaluation of intrusion detection systems. Mr. Porras states that SRI used eXpert to successfully detect several attacks in Lincoln Labs test data while stating “didn’t produce any meaningful results.” Mr. Porras’s assessment of eXpert on November 9, 1998 confirms my determination, described below, that as of November 9, 1998 SRI had implemented a fully functional version of eXpert.

November 9, 1998, SRI had a fully-functional version of eXpert that was integrated with P-BEST.

386. Based on my analysis of the source code placed into escrow by SRI, I have determined that prior to November 9, 1998 SRI had developed a fully-functional version of eXpert. As described below, prior to November 9, 1998, SRI had an implementation of P-BEST, code that integrated P-BEST with the rest of eXpert, and a complete suite of P-BEST signatures to analyze network traffic data.

a. P-BEST inference engine and translator

387. As mentioned earlier, the 1999 paper on P-BEST authored by Mr. Porras and Mr. Lindqvist states that SRI used P-BEST in its earlier IDES and NIDES intrusion detection systems well before 1998. Nevertheless, I reviewed several source code files, listed in Exhibit Y, underlying the P-BEST inference engine and P-BEST translator to confirm whether SRI had a working implementation of P-BEST as of November 9, 1998. My review of the files in Exhibit Y confirmed that as of November 9, 1998, SRI had a fully-functional implementation of the P-BEST software described in the 1999 P-BEST paper.

b. Integration between P-BEST and other parts of eXpert

388. I also analyzed the eXpert-related source code files to confirm that eXpert used P-BEST as its expert system. I have confirmed that as of November 9, 1998 eXpert used P-BEST as its expert system. The files that illustrate this integration are listed in Exhibit Z. It is important to note that while earlier versions of eXpert.c contained the relevant integration code, during the summer of 1998 that integration code was migrated to the source code file event.c.¹⁷²

¹⁷² See SYM_P_0549715 – 18.

c. Suite of P-BEST signatures for analyzing network traffic data

389. In order to determine the capabilities of eXpert as of November 9, 1998, I analyzed the P-BEST signatures that SRI had written as of November 9, 1998. The files that I analyzed are listed in Exhibit AA. Taken together, the rules in these files make up an entire suite of signatures for analyzing several types of network traffic data, including but not limited to data transfers, network packet data transfer commands, network packet data transfer errors, network connections, network connection requests, network connection acknowledgements, and ICMP packets that may include error codes.

i. The suite of P-BEST signatures analyzed several different types of network traffic data

390. I have analyzed the set of P-BEST rules contained in the file named ftp.pbest labeled as REQUEST # 22. These files analyze a number of different aspects of FTP-related traffic. FTP is an internet protocol that stands for "File Transfer Protocol." I am familiar with FTP and other application-layer network protocols such as HTTP, NFS, SMTP, and Telnet, and I may testify at trial regarding these and other application-layer network protocols. I have determined that the rules in this file monitor, among other things, data transfers, network packet data transfer commands, and network packet data transfer errors.¹⁷³

391. I have analyzed the set of P-BEST rules contained in the file named syn_attack.pbest and labeled REQUEST # 28. I am familiar with SYN attacks, and I may testify at trial regarding how SYN attacks operate and how they can be detected and

¹⁷³ My analysis has been confirmed in part by the testimony of Mr. Lindqvist, who modified parts of eXpert during the summer of 1998. Specifically, Mr. Lindqvist testified that the signatures in Request 22 analyze requests for data transfer. Lindqvist Depo. Tr. 110:3-11.

thwarted. I have determined that the set of rules in syn_attack.pbest analyze, among other things, network connections, network connection requests, and network connection acknowledgements.¹⁷⁴

392. I have analyzed the set of P-BEST rules contained in the file named rules.pbest and labeled Request # 30. The rules in this file analyzed network traffic data in order to detect an IP sweep. An IP sweep is a method by which to scan a network in order to ascertain the IP addresses of the nodes attached to the network. I am familiar with several types of sweeps and scans, and at trial I may testify as to how various sweeps and scans operate, how they can be detected, and how they can be thwarted. I have determined that the set of rules in Request # 30 analyze, among other things, ICMP packets that may include error codes.¹⁷⁵

ii. The P-BEST signatures were fully functional

393. I understand that SRI witnesses have characterized this suite of P-BEST rules aimed at analyzing network traffic as “experimental.”¹⁷⁶ I disagree. Based on my analysis of each of the rule sets comprising the suite of rules, I have concluded that each rule set would have been successful in detecting suspicious network activity. In fact, virtually identical rules from this suite ultimately appeared in a paper published by Mr. Porras and Mr. Lindqvist in 1999. Drafts of this paper from October 1998 show that by October 1998 Mr. Porras viewed several of these rules as mature and capable of detecting

¹⁷⁴ My analysis has been confirmed in part by the testimony of Mr. Lindqvist, who testified that the signatures in the file syn_attack.pbest (REQUEST # 28) analyzed “one type of network connection request” and “one type of network connection acknowledgment or lack thereof.” Lindqvist Depo. Tr. 119:16-120:1.

¹⁷⁵ My analysis has been confirmed in part by the testimony of Mr. Lindqvist, who testified that the signatures in the file marked as REQUEST # 30 may analyze ICMP error codes assuming that the ICMP_code field was referenced and assuming that the ICMP_code field corresponds to ICMP error codes. U. Lindqvist Depo. Tr. 137:6-17.

¹⁷⁶ Porras 30(b)(6) Depo. Tr. 190:5-23, 212:25-216:11.

SYN attacks.¹⁷⁷ In October 1998 Mr. Porras believed that the SYN attack eXpert rules were “done.”¹⁷⁸ Mr. Porras also wrote, “eXpert is already capable of detecting SYN Floods.”¹⁷⁹ Furthermore, SRI used eXpert during an important intrusion detection evaluation effort by their sponsor that occurred in late October and early November 1998.¹⁸⁰

4. eXpert was SRI’s best mode for practicing signature detection upon network traffic data

394. As noted previously, based upon my review of the source code, deposition testimony, and communications from the inventors described in this section, it is my opinion that the inventors also considered eXpert/P-Best to be their best mode of practicing “a signature matching detection method” as claimed in the ’212 patent. In fact, as of November 9, 1998, eXpert was the only signature engine among the EMERALD-related components that had been implemented.

5. eXpert was SRI’s best mode for analyzing network traffic data

395. I have compared the functionalities of estat and eXpert. It is my opinion that eXpert would have outperformed estat in analyzing network traffic data. Specifically, eXpert would have been able to detect more suspicious network activity than estat while at the same time generating less “false positives” than estat. In the context of intrusion detection systems, a false positive occurs when the system reports an intrusion when in fact there is no intrusion. A problem endemic to profile-based anomaly detection systems like estat is that they ordinarily generate more false positives than signature detection systems. Although I have not tested the estat component, I have

¹⁷⁷ See SRI 151720-40.

¹⁷⁸ SRI 277983-8002 at 277986.

¹⁷⁹ SRI 277983-8002 at 277987.

¹⁸⁰ See SRIE 0275123; SRI 0010873.

analyzed the statistical algorithms used by estat. Based on my analysis of those algorithms, I have concluded that estat would suffer the same problems regarding generating false positives that other profile-based anomaly detection engines suffer from.

396. In order to determine whether at the time of filing the inventors believed that eXpert was the best mode for analyzing network traffic data, I reviewed a twenty page document written by Mr. Porras in mid-October 1998 entitled, “EMERALD eXpert/estat—TCP/UDP/ICMP Analysis Summary.”¹⁸¹ In this document, Mr. Porras discusses 35 different attacks and analyzes whether estat, eXpert, or both estat and eXpert are well-suited to detect each attack. Exhibit BB summarizes Mr. Porras’s October 1998 conclusions. The October 1998 document clearly shows that Mr. Porras believed eXpert to be more effective overall than estat in detecting intrusions. As shown in Exhibit BB, Mr. Porras estimated that estat would detect eight of the 37 different attacks, while eXpert would detect 25 of the 37 different attacks. Furthermore, for several attacks, Mr. Porras wrote that estat would not be effective because “estat probably won’t see any significant anomalies here.”¹⁸²

6. In the alternative, the combination of eXpert and estat was the inventors’ best mode for analyzing network traffic data

397. Although it is my opinion that the inventors believed at the time of filing the application underlying the patents-in-suit that eXpert was the best mode for detecting “suspicious network activity,” I have been asked to assume for the sake of argument that the inventors believed that estat was better than eXpert for this purpose. If the inventors believed estat to be better than eXpert, it is my opinion that the inventors would have also viewed a combination of estat/eXpert to be the best mode for detecting suspicious

¹⁸¹ [SRI 277983].

¹⁸² See, e.g., SRI 27798328002 at 277986.

network activity; better than estat alone. In other words, even if the inventors believed that estat was better than eXpert in detecting suspicious network activity, the inventors would have believed that the best mode for detecting suspicious network activity would have been a combination of estat and eXpert. Thus, in my opinion the failure to disclose eXpert would still run afoul of the best mode requirement even if the inventors believed estat to be better at detecting suspicious network activity than eXpert.

7. eXpert was not disclosed in the patents-in-suit

398. Given my determination that eXpert was SRI's best mode for detecting suspicious network activity at the time of the filing of the application underlying the patents-in-suit, the next question in the best mode analysis is whether eXpert was disclosed in sufficient detail to allow one skilled in the art to practice it. In order to answer this question, I analyzed the patents-in-suit, the articles incorporated by reference into the patents-in-suit, and the source code appendix to the patents-in-suit.

399. I have reviewed the source code to appendix to the patents-in-suit and I have attempted to identify the names of each file in the appendix by comparing the files in the appendix to the files on the source code computer that SRI placed into escrow. Because the appendix was not included on the source code computer, it was not possible to do a character-by-character comparison of the files in the appendix and the files on the source code computer. In spite of this technical limitation, I have determined to the best of my ability the names and file paths for each of the files contained in the source code appendix. These file names and file paths are listed in Exhibit CC.

400. The source code files for the most substantial and important subcomponents of eXpert were not included in the source code appendix. First, the appendix contains none of the source code files—listed in Exhibit Y—for either the P-BEST inference engine or P-BEST translator.